arXiv:1611.00484v3 [cs.DB] 17 Nov 2016

RECOME: a New Density-Based Clustering Algorithm Using Relative KNN Kernel Density

Yangli-ao Geng[1] Qingyong Li[1] Rong Zheng[2] Fuzhen Zhuang[3] Ruisi He[1]

***Abstract*—Discovering clusters from a dataset with different shapes, density, and scales is a known challenging problem in data clustering. In this paper, we propose the RElative COre MErge (RECOME) clustering algorithm. The core of RECOME is a novel density measure, i.e., Relative K nearest Neighbor Kernel Density (RNKD). RECOME identifies core objects with unit RNKD, and partitions non-core objects into atom clusters by successively following higher-density neighbor relations toward core objects. Core objects and their corresponding atom clusters are then merged through α-reachable paths on a KNN graph. Furthermore, we discover that the number of clusters computed by RECOME is a step function of the α parameter with jump discontinuity on a small collection of values. A jump discontinuity discovery (JDD) method is proposed using a variant of the Dijkstra's algorithm. RECOME is evaluated on three synthetic datasets and six real datasets. Experimental results indicate that RECOME is able to discover clusters with different shapes, density and scales. It achieves better clustering results than established density-based clustering methods on real datasets. Moreover, JDD is shown to be effective to extract the jump discontinuity set of parameter α for all tested dataset, which can ease the task of data exploration and parameter tuning.**

I. Introduction

Clustering, also known as unsupervised learning, is a process of discovery and exploration for investigating inherent and hidden structures within a large data set [6]. It has been extensively applied to a variety of tasks such as information retrieval, text mining, Web analytics, bioinformatics, image processing [7], [8], [10], [16], [17], [21]–[24].

Many clustering algorithms have been proposed in different scientific disciplines [8], and these methods often differ in the selection of objective functions, probabilistic models, or heuristics adopted. K-means clustering is one of the most popular and simplest clustering algorithms, though it was proposed over 50 years ago [14]. K-means clustering iteratively assigns objects to K clusters with the nearest mean (or center) and updates the cluster centers. K-means clustering requires the knowledge of the number of clusters K, which is non-trivial problem in most real world applications. In addition, K-means often fails to discover non-spherical clusters [8].

Density based clustering methods can handle data with irregular shapes and discover clusters with arbitrary shapes. DBSCAN is a well-known density based clustering algorithm [5]. DBSCAN groups objects that are closely packed together, and marks outliers that lie in low-density regions. Recently, many improved methods have been proposed based on DBSCAN [13], [15]. OPTICS is another well-known density based clustering algorithm [2]. Its basic idea is similar to DBSCAN, but it addresses the problem of detecting meaningful clusters in data of varying density. Both DBSCAN and OPTICS can discover clusters with arbitrary shapes without the prior knowledge about the number of clusters K. They, however, face the challenge of choosing appropriate parameter values. Recently, a novel density based clustering method was proposed by fast search-and-find of density peaks (FDP) [18]. This algorithm assumes that cluster centers are surrounded by neighbors with lower local density and that they are at a relatively large distance from any point with higher density. Though FDP is efficient and effective, it lacks any efficient quantitative criterion for judging cluster centers. Accordingly, some approaches have been proposed to improve FDP [3], [11].

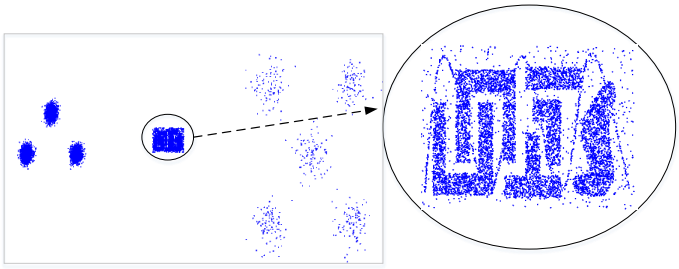

Fig. 1: A synthetic dataset with complex distribution, which includes clusters with different shapes, density, and scale.

Density-based clustering methods have the advantage of discovering clusters with arbitrary shapes and dealing with noisy data, but they face two challenges. First, traditional density measures are not adaptive to datasets with complex distribution. For example, the dataset in Fig. 1 includes clusters with different shapes, density, and scales. The density defined in [5], [18] would not be applicable. Second, performance of traditional methods (e.g., DBSCAN and FDP) is sensitive to parameter selection, and it is non-trivial to set these parameters properly for different datasets.

Aiming to address these challenges of density based clustering methods, we presents the RElative COre MErge (RECOME) clustering algorithm, which is based on the novel density measure, i.e., Relative K nearest Neighbor Kernel Density (RNKD). RECOME firstly identifies *core objects* corresponding to data points with RNKD equal 1. A core object

[1]School of Computer and Information Technology, Beijing Jiaotong University, Beijing, China. {gengyla,liqy,ruisi.he}@bjtu.edu.cn

[2]Department of Computing and Software, McMaster University, Hamilton, Canada. rzheng@mcmaster.ca

[3]Key Lab of Intelligent Information Processing of Chinese Academy of Sciences, Beijing, China. zhuangfz@ics.ict.ac.cn

and its descendants, which are defined by a directed relation (i.e., *higher density nearest-neighbor*), form an *atom cluster*. These atom clusters are then merged using a novel notion of α-connectivity on a KNN graph. RECOME have been evaluated using three synthetic datasets and six real datasets, and the experiment results demonstrate that RECOME outperforms DBSCAN, KNNCLUST, and FDP. Furthermore, we discover that the clustering results of RECOME can be characterized by a step function of its parameter α, and therefore put forward an effective jump discontinuity discovery (JDD) algorithm to extract all but finite jump discontinuity values. In summary, this work makes the following contributions.

1) The new density measure, RNKD, is more robust to variability in density compared with absolute density measures in DBSCAN and non-parametric kernel density. In addition, RNKD is instrumental in finding out clusters with relative low density.
2) RECOME can avoid the "decision graph fraud" problem [11] of FDP and can handle clusters with heterogeneous density. Furthermore, RECOME has linear computational complexity if the distance matrix between objects is computed in advance.
3) JDD can extract all jump discontinuity of parameter α for any dataset. It will greatly benefit parameter selection in real-world applications.

This paper is organized as follows. Section II introduces the related work. Section III presents the new density measure RNKD, and Section IV describes the proposed clustering method RECOME. Section V analyzes the parameter α of RECOME and presents the auxiliary algorithm JDD. Section VI demonstrates experimental results. Finally, we conclude the paper in Section VII.

II. Related work and notations

In this section, we review some related work on density-based clustering methods.

DBSCAN [5] is one of the most classic density-based methods. In DBSCAN, the *cut-off density* of an object o is defined as the number of objects falling inside a ball of radius ϵ centred at o. If the cut-off density of o is higher than $MinPts$, o is regarded as a key object. When the distance between two key objects is less than ϵ, they are called density-reachable. Density-reachable key objects form basic clusters. A non-key object is assigned to a basic cluster if it is within ϵ distance to a key object in the respective cluster; otherwise, the non-key object is treated as noise. In DBSCAN, ϵ and $MinPts$ need to be carefully tuned as they would affect the number and formations of the final clusters. Furthermore, DBSCAN cannot handle clusters of heterogeneous densities well as both parameters, even properly selected, are fixed. In contrast, in this work, we introduce the novel concept of RNKD, which in effect homogenizes density measures and thus can handle clusters of different densities.

Kernel density is a well-known alternative to cut-off density. Kernel density is continuous and is less sensitive to parameter selection. However, traditional kernel density methods face the same difficulty as DBSCAN when dealing with complex datasets that contain clusters with different densities [4], and tend to confuse low-density clusters with noise. Consequently, KNN kernel density has been applied to handle such situations [20]. The proposed RNKD estimation is inspired by the KNN kernel density with further improvement allowing the inclusion of low-density clusters.

Rodriguez and Laio proposed a novel density-based clustering method by finding density peaks [18]. FDP discovers clusters by a two-phase process. First, a decision group method is devised to determine cluster centers, called *density peaks*. In this step, a local density measure is computed for each data point based with the number of data points in its d_c neighborhood. Second, remaining objects are assigned to the same cluster as its nearest neighbor with higher density. FDP is effective in finding clusters with different shapes and densities. However, as mentioned by the authors, the cluster centers and consequently the number of clusters are sensitive to the choice of the cut-off parameter d_c in computing the local density. In contrast, RECOME adopts an agglomerative procedure to merge atom clusters (analogous to those resulting from "local peaks"). We have further investigated the JDD algorithm to allow users to assess the parameter α in RECOME.

Karypis *et al.* [9] proposed a hierarchical clustering method based on KNN graphs, where each object has an edge to each of its K nearest neighbors. The clustering method applies a graph partitioning algorithm to divide the KNN graph into a large number of relatively small sub-clusters with the objective of minimizing edge cut. Sub-clusters are then merged based on their similarity characterized by relative interconnectivity and relative closeness. The KNN graph in [9] is used for initial partitions (equivalence of atom clusters in RECOME). In RECOME, KNN graphs are utilized to determine the final clusters from atom clusters.

Key notations in the paper are listed in Table I.

TABLE I: Main notations used throughout the paper

Notation	Description
$\lvert\cdot\rvert$	Absolute value of a scalar or cardinality of a set
V	Dataset, i.e., $V = \{v_1, v_2, \ldots, v_n\}$ where $v_i \in \mathbb{R}^m$, m is the feature dimension of objects. Lower case symbols u, v, w, u_i, v_i, w_i denote elements of V
$N_K(u)$	The K nearest neighbor set of u. $N_K(u) = \{u_1, u_2, \ldots, u_K\}$, where u_i is the i^{th} nearest neighbor of u, and $u_i \triangleq N_i(u) \backslash N_{i-1}(u)$
$d(u, v)$	Euclidean distance between u and v

III. Relative KNN Kernel Density Estimation

In this section, we introduce the proposed density measure, RNKD, which is a cornerstone of RECOME.

Density estimation is crucial in density-based clustering algorithms, and different density forms have been adopted in literature. For instance, the cut-off density in DBSCAN is defined as the number of objects in an ϵ-ball centred at the respective data point. As discussed in Section II, the parameter ϵ needs to be carefully tuned based on the characteristics

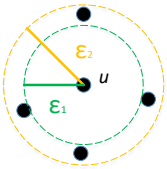

(a) ϵ cutoff density

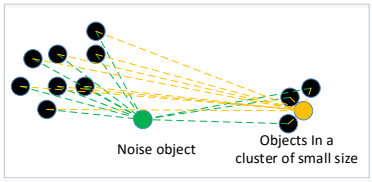

(b) Kernel density

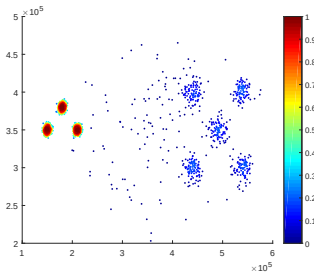

(c) NKD

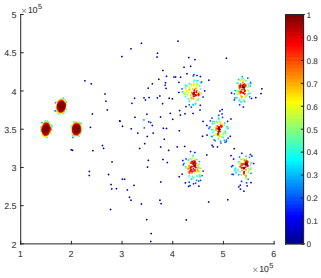

(d) RNKD

Fig. 2: Comparison of different density measures. (a) Sensitivity to ϵ in DBSCAN; (b) Non-parametric kernel density methods are sensitive to noise; (c) NKD estimates may mistake low-density clusters as noise; (d) RNKD allows discrimination of low-density clusters from noise. The temperature of data points in (c) and (d) indicates the (normalized) NKD and RNKD, respectively.

of the target dataset for good clustering performance, and a fixed setting is not suitable for inhomogeneous densities. An illustrative example is given in Fig. 2a, which demonstrates that a small variation in ϵ can result in drastic differences in density estimation. On the other hand, non-parametric kernel density estimation methods such as that in [20], though superior in handling complex distributions, are sensitive to noise as they consider all points in the dataset (Fig. 2b).

Our proposed RNKD estimation is inspired by KNN kernel density [20], which only considers the objects in $N_K(u)$ for the density estimation of object u. The K-nearest neighbor kernel density (NKD) of object u, denoted by $\rho_K(u)$, is defined as,

$$\rho_K(u) = \sum_{v \in N_K(u)} \exp\left(-\frac{\mathrm{d}(u,v)}{\sigma}\right), \qquad (1)$$

where $\sigma = \frac{\sum_{v \in V} \mathrm{d}(v, v_K)}{|V|}$, namely the mean of the distance between v and its Kth nearest neighbor for all v in V.

NKD allows easy discrimination of outliers, but may mistake low-density clusters as outliers. One such example is shown in Fig. 2c, where low-density clusters on the right may be mistaken as noise. The key observation is that though the (absolute) NKD of a low-density cluster is low compared to the high-density clusters, its NKD is still comparatively high in its surroundings. This observation motivates us to introduce a new density measure called RNKD as follows. The definition of RNKD will be supplied in the next version. Thus, RNKD in effect homogenizes the density estimation across clusters of different densities. Figure 2 shows that both dense clusters and sparse clusters have high RNKD values.

Note that in computing NKD in (1), a Laplace kernel is adopted as opposed to commonly used Gaussian kernels. This is motivated by two considerations, i.e., i) both types of kernels lead to similar performance, and ii) Laplace kernel is more efficient to compute.

Choosing K

The choice of K is expected to affect the quality of density estimates. In particular, when K is large, in the extreme case when $K = |V|$, RNKD degenerates to a global density measure. The resulting clustering algorithm may disregard low-density clusters. Furthermore, a large K increases the computation complexity in determining RNKD. On the other hand, when K is small, say $K = 1$, the resulting density estimation is determined only by pair-wise distance and is likely to have high fluctuation.

Through experimental study, we observe that choosing K in the range of $\left[\sqrt{|V|}, 2\sqrt{|V|}\right]$ will provide stable and good performance. Further empirical evidence will be provided in Section VI. As part of the future work, we plan to investigate the theoretical basis for the choice of K.

IV. RECOME Clustering Algorithm

Now we are in the position to present the RECOME algorithm based on RNKD. The basic idea is as follows. First, we identify *core objects* corresponding to data points of peak relative density. These core objects serve as centers of sub-clusters, called *atom clusters*, which will be further merged through connected paths on a KNN graph. Thus, RECOME is in essence an agglomerative hierarchical clustering method. This section will supplied in the next version.

V. Jump Discontinuity Discovery for Parameter α

As evident from the description of RECOME, parameter α is crucial to the clustering result for a fixed K. Take the example in Fig. 3. Starting from the same set of core objects and atom clusters, different values of α may result in different numbers of clusters. In particular, as α increases, cluster granularity (i.e., the volume of clusters) decreases and cluster purity increases. Thus, a pertinent question is how to select a proper α.

Parameter selection to control the granularity and the purity of clusters is a fundamental problem in clustering algorithms. Objective metrics of clustering quality such as F-measure and normalized mutual information (NMI) are only applicable when the ground truth is available. In many application scenarios such as community detection in social networks [16], computer vision [24], ground truth information is not readily available. However, it is often desirable to provide tuning knobs to end users so that the users can assess the clustering outcomes from different parameter settings based on domain knowledge.

A. Jump Discontinuity Set

In this section, we analyze the impact of α on clustering results. In particular, we find that though α varies continuously in the interval $[0, 1]$, the number of clusters computed by RECOME is a step function of α. For instance, in Fig. 3, only five clustering outcomes are possible for the example dataset, corresponding to α in the ranges of $[0, 0.6)$, [0.6, 0.7),[0.7, 0.8), $[0.8, 0.9]$ and $[0.9, 1)$. The numbers of resulting clusters are 1, 3, 5, 7 and 8, respectively. It is desirable to have a small collection of α values (or ranges) that affect the clustering result as the processes of parameter tuning by developers or parameter selection by domain experts can be simplified.

We formalize the above intuition by first introducing the notion of jump discontinuity set.

Definition 1. Given a data set V and an input parameter K, an ascending list $L = \{\alpha_1, \alpha_2, \ldots, \alpha_l\}$ is called a **jump discontinuity** (JD) set if the number of resulting clusters from RECOME, $\#(V, K, \alpha)$ is a step function of α with jump discontinuity at $\alpha_1, \alpha_2, \ldots, \alpha_l \in [0, 1]$ from left to right.

It is easy to see that $\#(V, K, \alpha)$ is a non-decreasing function of α. By definition, each JD in L yields a unique clustering result. From all the JDs in L, we can produce all possible clusters using RECOME. Recall that O is the set of core objects and $|O|$ is the maximum number of clusters attainable by RECOME. Trivially, $|L| \leq |O|$.

B. JD Discovery Algorithm

Without loss of generality, we assume that the KNN graph G is connected. If G is disconnected, each connected component can be treated individually. The basic idea of the JD discovery algorithm is to first find the maximum capacity paths among core components and then compute the JD set from the weights of the "weakest" nodes on these paths.

To facilitate the discussion, we first introduce some terms and discuss their properties. Consider a path $p = \langle u, w_1, w_2, \ldots, w_s \equiv v \rangle$, Its *left-open path capacity* is defined as $c(p) = \min_i \rho^*_K(w_i)$. Denote $\omega(p) = \arg\min_i \rho^*_K(w_i)$ as *the weakest point* on path p. Suppose there are l paths from u to v, denoted by $p_1, p_2, \ldots, p_l$. The *left-open inter-node capacity* between u and v is defined as $c(u, v) = max_l c(p_l)$, and *the maximum left-open capacity path* from u to v is denoted by $p_{max}(u, v) = \arg\max_l c(p_l)$. In other words, v is α-reachable from u for all $\alpha < c(u, v)$. For the ease of presentation, we drop the term "left-open" from here on unless otherwise noted. The following properties immediately follow from the definition.

Lemma 1. *Given a path* $p = \langle u, w_1, w_2, \ldots, w_s \equiv v \rangle$, *consider three sub-segments* $p_1 = \langle u, w_1, w_2, \ldots, w_{s_1} \rangle$, $p_2 = \langle w_{s_1}, w_{s_1+1}, \ldots, w_s \equiv v \rangle$, *and* $p_3 = \langle w_{s_1}, w_{s_1+1}, \ldots, w_{s_2} \rangle$, *where* $1 \leq s_1, s_2 \leq s$. *We have i)* $c(p) = min(c(p_1), c(p_2))$, *ii)* $c(p) \leq c(p_3)$, $\forall s_1, s_2, s.t., 1 \leq s_1, s_2 \leq s$.

Lemma 2 (Bellman-Ford Equation)**.** *The inter-node capacity between* u *and* v *satisfies:*

$$c(u, v) = \max_{(u,w) \in E} \left(\min(\rho^*_K(w), c(w, v))\right). \tag{2}$$

From Lemma 1 and Lemma 2, we observe that our definitions of path and inter-node capacities bare some similarity to least cost paths in graphs. In fact, we can use a variant of the

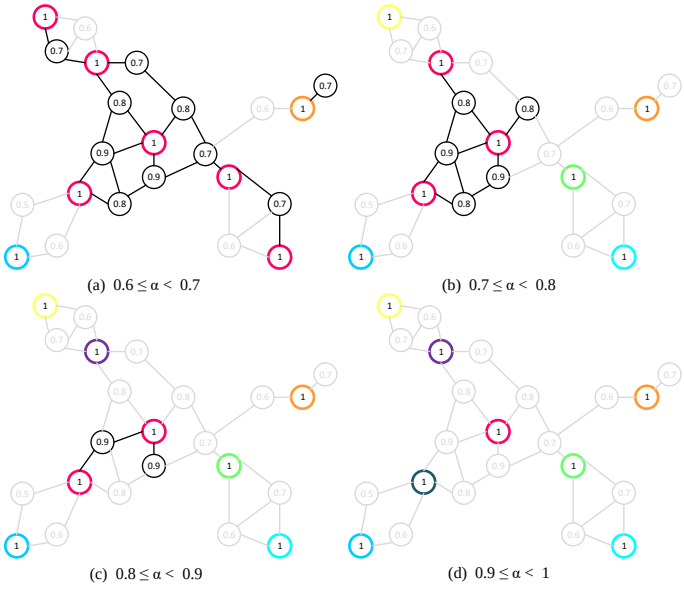

(a) $0.6 \leq \alpha < 0.7$ (b) $0.7 \leq \alpha < 0.8$ (c) $0.8 \leq \alpha < 0.9$ (d) $0.9 \leq \alpha < 1$

Fig. 3: Example of merging core objects results for different values of α. KNN graph is shown and each object (in circle) is marked with its relative density. Core objects of the same color belong to the same final cluster.

Dijkstra's algorithm to determine the maximum capacity paths from any core object (the source) to any other core objects.

Starting from a core component u as the initial node, we execute a variant of the Dijkstra's algorithm as follows.

1) Assign to every node $v \in V$ a tentative capacity value $\lambda(v)$: set it to 1 for the initial node and to 0 for all other nodes.
2) Set the initial node as current. Mark all other nodes unvisited. Create a set of all the unvisited nodes called the *unvisited set* $\mathcal{U}$, i.e., $\mathcal{U} = V \backslash \{u\}$.
3) For the current node v, consider all of its unvisited neighbors w and update their capacity as,

   $$\lambda(w) = \max(\lambda(w), \min(\lambda(v), \rho_K^*(w))), \forall w \in N_K(v) \cap \mathcal{U}$$

   Mark the current node as visited and remove it from the unvisited set, namely, $\mathcal{U} = \mathcal{U} \backslash \{v\}$.
4) If all core objects are marked as visited, terminate the process.
5) Otherwise, select the unvisited node with the highest tentative capacity, set it as the new current node, and go back to Step 3.

Similar to the Dijkstra's algorithm, the correctness of the afore-mentioned procedure hinges upon the following invariant property,

Lemma 3. *For each visited node $v \in V \backslash \mathcal{U}$, $\lambda(v)$ is the inter-node capacity from the source node u to v; and for each unvisited $v \in \mathcal{U}$, $\lambda(v)$ is the largest capacity via visited nodes only from source u to v.*

Proof: See Appendix. ∎

From Lemma 1, it is easy to see that the maximum capacity paths from a core object u to all other core objects forms a tree since loops do not increase the path capacity. Denote this tree by $M^*(u)$. To this end, we have found all the maximum capacity paths from source u to any $v \in O \backslash \{u\}$. Recall $\omega(p)$ is the weakest point of path p. For each maximum capacity path $p_{max}(u, v)$, we find its weakness point and include its associated relative density in the list L. In other words, $\forall v \in O \backslash \{u\}$, $L = L \cup \{\rho_K^*(\omega(p_{max}(u, v)))\}$.

From the definition of maximum capacity paths, it is easy to show that,

Lemma 4. *$\forall u, v \in O$ and $\alpha \in [0, 1]$, if $u \xrightarrow{\alpha} u$ in G, then $u \xrightarrow{\alpha} v$ in $M^*(u)$.*

Repeating the process for each core object, we will obtain the final JD list by removing duplicates and ordering the elements in the ascending order.

The parameter JD discovery algorithm is thus summarized in Algorithm 1. The steps are illustrated in Fig. 4.

Theorem 1 presents the correctness of Algorithm 1 in finding all JD parameters.

Theorem 1. *Let $L = \{\alpha_1, \alpha_2, \ldots, \alpha_l\}$ be the list produced by Algorithm 1. Then,*

- *(Correctness) All elements in L produced by Algorithm 1 are JD parameters.*

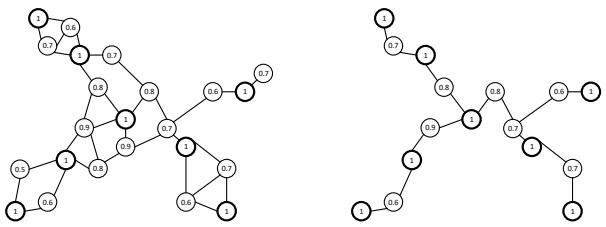

(a) G (b) M^*

Fig. 4: An illustration of KNN graph G and a maximum capacity tree M^* from one core object. In this example, the maximum capacity tree with each core object as the source node coincides with one another.

Algorithm 1: JD Discovery for parameter of α

Input: KNN graph G and core object set O
Output: Sorted list JD list L

```
L ← Φ;
for u ∈ O do
    Find maximum capacity paths from u to v ∈ O\{u};
    for v ∈ O\{u} do
        L ← L ∪ {ρ*_K(ω(p_max(u, v)))};
    end
end
```

- *(Completeness) If $\alpha \notin L$ and $\alpha \in [0, 1]$, it is not a JD parameter.*

Proof: See Appendix. ∎

VI. Experimental Results and Analysis

In this section, we test the proposed RECOME over three 2-D synthetic datasets and compare with three representative algorithms over six real datasets. All experiments are implemented in Microsoft Visual C++ 2015 14.0.24720.00 on a workstation (Windows 64 bit, 4 Intel 3.2 GHz processors, 4 GB of RAM).

A. Experiment Setup

1) Dataset: Three 2-D synthetic datasets and six real datasets from the UCI repository [12] are used in this work. The three synthetic datasets are described as follows. *Dataset1* is an unbalance datset with sparse noise, which contains 8 classes of different density and size, 6617 points in total, and 117 noise points; *dataset2* refers to the t4.8k dataset, which includes 6 classes of nonconvex shape and 8000 points; *dataset3* represents the mixture of t4.8k dataset and Unbalance dataset and contains 14 classes of different shapes, density and scale.

The six real datasets are the *Iris* dataset, the *wine* dataset, the *seeds* dataset, the page blocks classification dataset (*PBC*), the land-sat satellite dataset (*LS*), and the pen-based recognition of handwritten digits dataset (*PBRHD*). We summarize the characteristics of the six datasets in Table II.

TABLE II: Summary of the six real datasets

Dataset	Number of clusters	Number of features	Number of objects
Iris	3	4	150 (50,50,50)
Seeds	3	7	210 (70,70,70)
Wine	3	13	178 (59,71,48)
PBC	5	10	5473 (4913,329,28,88,115)
LS	6	36	6435 (1533,703,1358,626,707,1508)
PBRHD	9	16	10992(1143,1144,1055,1144,1055,1056,1142,1055,1055)

2) Baseline Methods and Settings: The algorithms used for comparison are listed as follows.

DBSCAN [5]: We estimate the density by Gaussian kernel instead of cut-off density because it produces better performance than other density measures. The density of object u is calculated as

$$\rho(u) = \sum_{v \in V - \{u\}} \exp\left(\frac{\mathrm{d}(u,v)}{\sigma}\right)^2,$$

where σ equals the value of the top β percent distance among all object pairs. We experiment with different β values in $\{1,2,3,4,5,6,7,8,9,10\}$. The parameter $MinPts$ is determined by $\lambda \sum \rho(v)/|V|$, where λ is chosen from $\{0.25, 0.5, 0.75, 1, 1.25, 1.5\}$.

KNNCLUST [20]: According to the recommendation in [20], parameter K is taken from $[|V|/4C, |V|/2C]$ with a step size of 10% of the range, where C is the real class number.

FDP [18]: FDP calculates density the same way as DBSCAN. It take the top large C objects as cluster centers according to γ values, where C is the real class number.

For RECOME, we set parameter K with the fixed value $\sqrt{|V|}$. As for parameter α, we enumerate all JD values extracted by JDD algorithm and select the best one.

The best performance are reported for the baseline methods by searching through the afore-mentioned parameter space.

3) Clustering Evaluation: For the 2-D synthetic datasets, we visualize clustering results. For real datasets, two measures are calculated based on the ground truth of the datasets. One is the normalized mutual information (NMI) [19]. The other one is the F value, $F = \frac{2 \times PB \times RB}{PB \times RB}$, where PB and RB refer to precision b-cubed and recall b-cubed [6], respectively. We use this measure because PB and RB have been shown superior than other indexes [1]. Given the dataset $V = \{v_1, v_2, \ldots, v_n\}$, the actual class labels $\{l_1, l_2, \ldots, l_n\}$, and cluster labels $\{c_1, c_2, \ldots, c_n\}$, we define

$$\text{Correctness}(v_i, v_j) = \begin{cases} 1 & \text{if } l_i = l_j \Leftrightarrow c_i = c_j \\ 0 & \text{otherwise} \end{cases}.$$

Thus, PB and RB are computed as follows:

$$PB = \frac{1}{n} \sum_{i=1}^{n} \frac{\sum_{i \neq j, c_i = c_j} \text{Correctness}(v_i, v_j)}{|\{v_j | i \neq j, c_i = c_j\}|},$$

$$RB = \frac{1}{n} \sum_{i=1}^{n} \frac{\sum_{i \neq j, l_i = l_j} \text{Correctness}(v_i, v_j)}{|\{v_j | i \neq j, l_i = l_j\}|}.$$

Both NMI and F fall in $[0, 1]$, and a higher value denotes better clustering performance.

B. Experimental Results

1) Results for Synthetic Dataset: Figure 5 demonstrates the clustering results of the four methods for the three datasets. DBSCAN divides the dataset2 (the second column) well, but fails to discover the sparse clusters in dataset1 (the first column) and dataset3 (the third column). FDP has inferior performance with dataset2 and dataset3 because their clusters have non-convex shape and are in different scales. FDP handles clusters with different density better than DBSCAN and KNNCLUST. KNNCLUST achieves the worst performance for the three datasets, because it is a partitioning method in principle and it calculates memberships of objects using KNN kernel estimation instead of the nearest neighbor rule. We observe that RECOME achieves superior clustering results for the datasets with different density, shapes, and scales.

These results indicate that RECOME can indeed handle datasets with arbitrary shapes. Furthermore, compared to other density based methods, the proposed density measure RNKD allows discovery of clusters with relative low density.

2) Results for Real Dataset: TABLE III lists the NMI and F values obtained by the four algorithms on the six datasets. RECOME has the best performance for Iris, Wine, and PBRHD; whereas its performance on other datasets is comparable with DBSCAN, KNNCLUST, and FDP. We have to note that KNNCLUST and FDP are given the true cluster number, which can benefit to improve their performance. The average NMI and average F of RECOME are greater than those of other three methods. Therefore, we observe that RECOME works well on real datasets and achieves better performance than DBSCAN, KDDCLUST, and FDP on average.

C. Parameter Analysis

First, we investigate the influence of parameter K, which determines the density measure RNKD and the topological structure of the KNN graph. In this experiment, K takes values from $[0, 5\sqrt{|V|}]$ with a step of $\sqrt{|V|}/10$. We first compute the JD set L for each fixed K, and then the best clustering results for NMI and F are reported. Figure 8 shows the clustering performance with respect to K. For clarity, we let $K = k\sqrt{|V|}/10$ and vary the index k instead K. It can be observed from Fig. 6 that RECOME can obtain stable

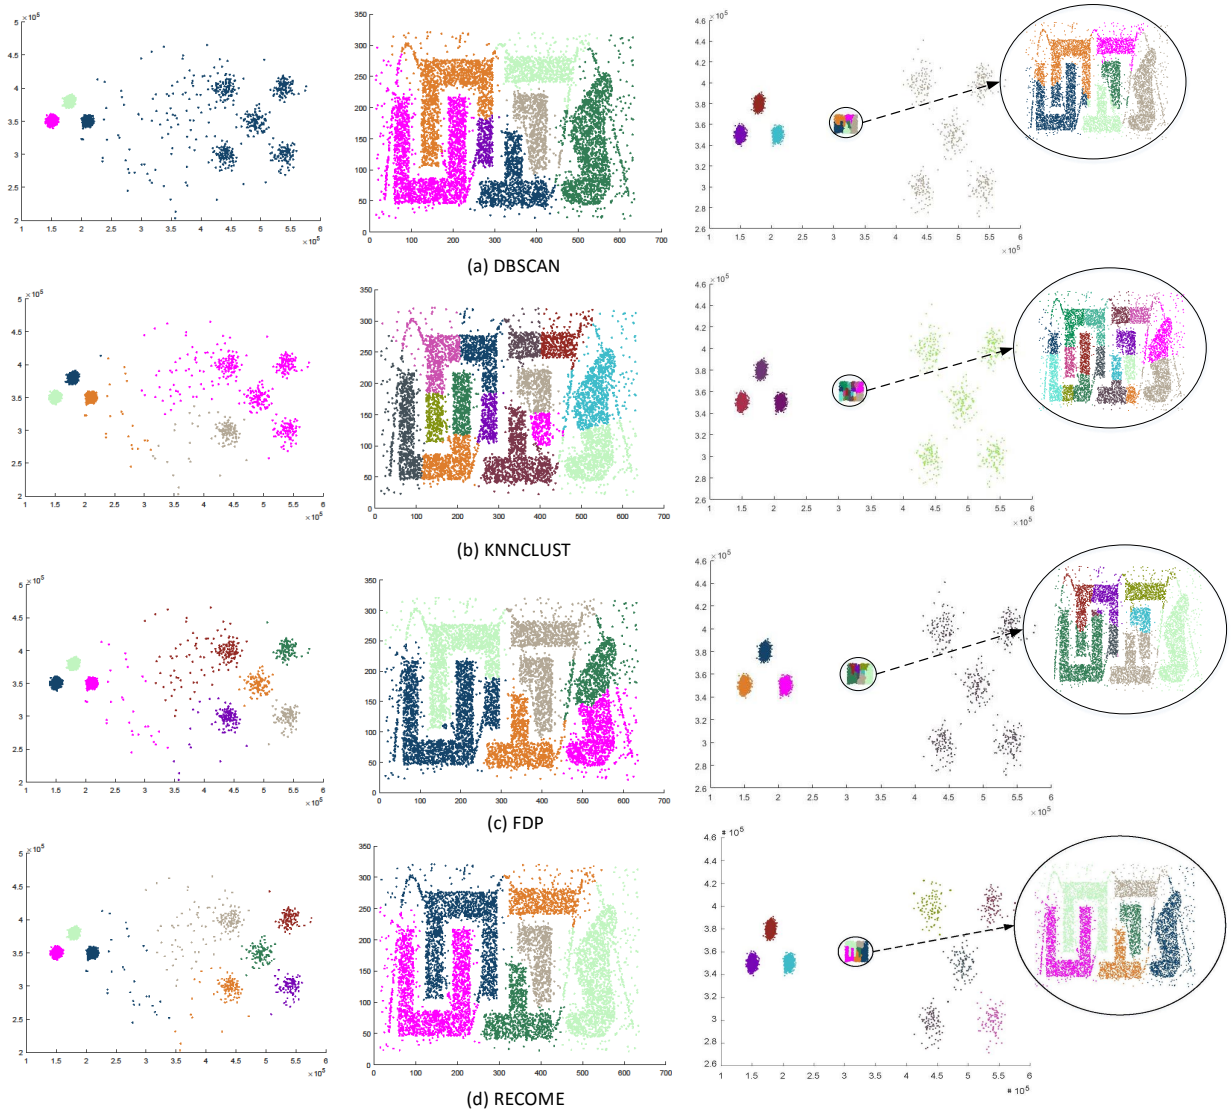

Fig. 5: Visualization of clustering results of DBSCAN, KNNCLUST, FDP, and RECOME for dataset1 (left column), dataset2 (middle column), and dataset3 (right column), respectively. Note that different clusters are displayed with different colors.

TABLE III: Performance comparison of the four methods on the six real datasets

	DBSCAN		KNNCLUST		FDP		RECOME ($K = \sqrt{\lvert V \rvert}$)		
Dataset	NMI	F	NMI	F	NMI	F	NMI	F	K
Iris	**0.771**	**0.827**	0.714	0.749	0.744	0.814	**0.771**	**0.827**	12
Seeds	0.747	0.858	0.674	0.784	**0.756**	**0.867**	0.714	0.848	14
Wine	0.630	0.757	0.712	0.808	0.726	0.823	**0.753**	**0.845**	13
PBC	0.273	**0.901**	0.154	0.895	**0.288**	0.856	0.273	0.897	73
LS	0.517	0.605	**0.653**	0.686	0.632	0.679	0.639	**0.691**	80
PBRHD	0.672	0.618	0.812	0.770	0.748	0.691	**0.836**	**0.785**	104
Average	0.602	0.761	0.620	0.782	0.649	0.788	**0.658**	**0.811**	

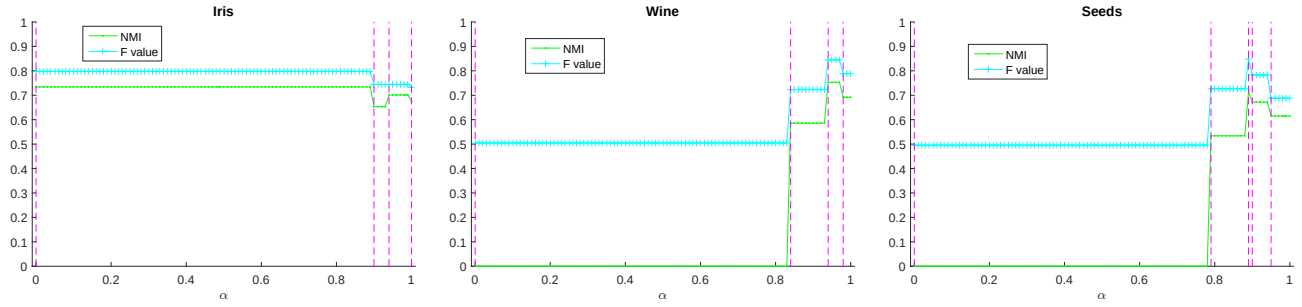

(a) L={0, 0.90, 0.94, 1}
$\mathcal{C}$={2, 3, 4, 6}

(b) L={0, 0.84, 0.94, 0.98}
$\mathcal{C}$={1, 2, 3, 4}

(c) L={0, 0.79, 0.89, 0.90, 0.95}
$\mathcal{C}$={1, 2, 3, 4, 5}

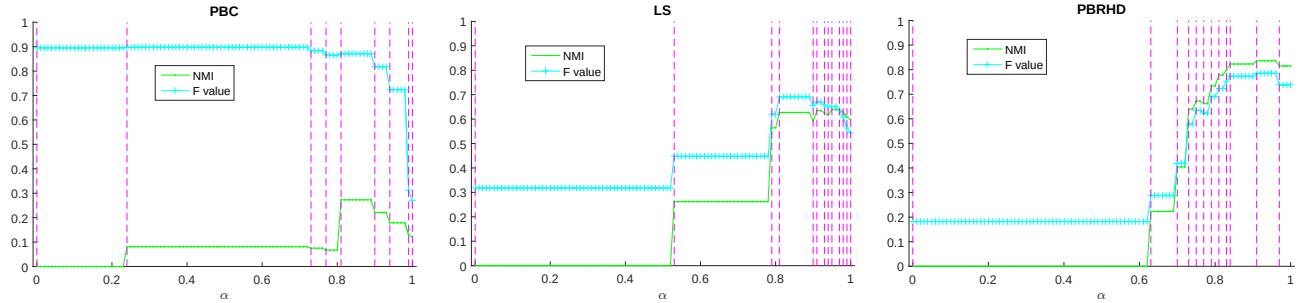

(d) L={0, 0.24, 0.73, 0.77, 0.81, 0.90, 0.94, 0.99, 1}
$\mathcal{C}$={1, 2, 3, 4, 5, 6, 8, 11, 19}

(e) L={0, 0.53, 0.79, 0.81, 0.90, 0.91, 0.93, 0.94, 0.95, 0.98, 0.99, 1}
$\mathcal{C}$={1, 2, 3, 4, 5, 7, 8, 9, 10, 14, 15, 18}

(f) L={0, 0.59, 0.60, 0.69, 0.71, 0.77, 0.78, 0.80, 0.81, 0.83, 0.85, 0.91, 0.97, 0.98}
$\mathcal{C}$={1, 2, 3, 4, 6, 7, 9, 12, 14, 16, 17, 18, 19}

Fig 9: The clustering performance as a function of α. The elements of JD set L produced by Algorithm 2 are marked by purple dashed line. The JD set L and corresponding collection of cluster numbers C are listed in the caption.

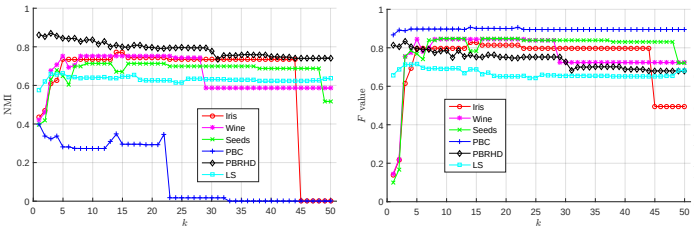

(a) NMI as a function of K

(b) F as a function of K

Fig. 6: The clustering performance as a function of K. The tick on x axis refers to the step index k, namely $K = k\sqrt{|V|}/10$.

performance when K ranges in $[\sqrt{|V|}, 2\sqrt{|V|}]$. Thus, we recommend to choose $K = \sqrt{|V|}$ as it achieves a good balance between clustering performance and computation complexity.

As discussed in Section V, clustering results of RECOME is a step function of α. In the next experiment, we fix K at $\sqrt{|V|}$ and search α in $[0, 1]$ with a step size 0.01. The clustering performance (i.e., NMI and F) as a function of α is plotted in Fig. 9. It can be observed that all curves are staircase though they have different monotonicity. In addition, each α yields a clustering result with distinctive cluster number C. From the set of results, it appears that setting α in the range of $[0.8, 0.9]$ can give good performance.

VII. Conclusion

In this paper, we have presented a new density based clustering method RECOME for discovering clusters with different shapes, density, and scales. Firstly, a novel density measure RNKD has been proposed to overcome the problems faced by cut-off density and the conventional KNN kernel density. Secondly, RECOME exploits RNKD to find core objects and corresponding atom clusters, and then merges the core objects based on the KNN graph efficiently. It could effectively balance cluster granularity and purity using a graph-theoretical construct. Finally, we discovered that the number of clusters obtained by RECOME is a step function of α, and proposed an auxiliary algorithm JDD to help end users to select the parameter α. Extensive experiment evaluation using both synthetic and real-world datasets has demonstrated the effectiveness of the proposed method.

VIII. Appendix

A. Proof of Lemma 3 in Section V

We prove by induction. The base case is when there is just one visited node, namely the initial node source, and the result trivially holds.

Assume the result is true for $m-1$ visited nodes. Now we choose an edge (v, w) where w has the largest $\lambda(w)$ of any unvisited node and the edge (v, w) is such that $\lambda(w) = min(\lambda(v), \rho^*_K(w))$. By Lemma 2, $\lambda(w)$ must be the

maximum capacity from source to w because if there were a path of higher capacity, and if w' was the first unvisited node on that path then by hypothesis $\lambda(w') > \lambda(w)$ creating a contradiction. Similarly if there was a higher capacity path to w without using unvisited nodes $\lambda(w)$ would have been less than $min(\lambda(v), \rho^*(w))$.

After processing w it will still be true that for each unvisited node w', $\lambda(w')$ is the maximum capacity from source to w' using visited nodes only, since if there were a higher capacity path which does not visit w we would have found it previously, and if there is a shorter path using w we update it when processing w.

B. Proof of Theorem 1

Due to space limit, we only provide a proof sketch.

To show correctness, starting from α_1, at step i, we remove all non-core nodes with weights less than or equal to α_i. Let the remaining graph be G_l. From the construction of L in Algorithm 1, there exists a pair of core objects u and v, such that $u \xrightarrow{\alpha_{i+1}} v$ in G. By Lemma 4, $u \xrightarrow{\alpha_{i+1}} v$ holds in G_l as well.

To show completeness, we prove by contradiction. If a JD parameter $\alpha \notin L$ exists, without lost of generality, $\alpha_i < \alpha < \alpha_{i+1}$. For each pair u and v that are connected in G_i, their inter-node capacity is no less than α_{i+1}. This contracts with the fact that α is a JD parameter.

REFERENCES

[1] Enrique Amig, Julio Gonzalo, Javier Artiles, and Felisa Verdejo. A comparison of extrinsic clustering evaluation metrics based on formal constraints. *Information Retrieval*, 12(4):461–486, 2009.

[2] Mihael Ankerst, Markus M Breunig, Hans-Peter Kriegel, and Jörg Sander. Optics: ordering points to identify the clustering structure. In *ACM SIGMOD*, volume 28, pages 49–60, 1999.

[3] Nurjahan Begum, Liudmila Ulanova, Jun Wang, and Eamonn Keogh. Accelerating dynamic time warping clustering with a novel admissible pruning strategy. In *ACM SIGKDD*, pages 49–58, 2015.

[4] Dorin Comaniciu and Peter Meer. Distribution free decomposition of multivariate data. *Pattern analysis & applications*, 2(1):22–30, 1999.

[5] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *ACM SIGKDD*, volume 96, pages 226–231, 1996.

[6] Jiawei Han, Micheline Kamber, and Jian Pei. *Data mining: concepts and techniques*. Elsevier, 2011.

[7] Juhua Hu, Qi Qian, Jian Pei, Rong Jin, and Shenghuo Zhu. Finding multiple stable clusterings. In *IEEE ICDM*, pages 171–180, 2015.

[8] Anil K Jain. Data clustering: 50 years beyond k-means. *Pattern recognition letters*, 31(8):651–666, 2010.

[9] George Karypis, Eui-Hong Han, and Vipin Kumar. Chameleon: Hierarchical clustering using dynamic modeling. *Computer*, 32(8):68–75, 1999.

[10] Ji Kubalk, Pavel Tich, Radek indel, and Raymond J. Staron. Clustering methods for agent distribution optimization. *IEEE Transactions on Systems Man & Cybernetics Part C*, 40(1):78–86, 2010.

[11] Zhou Liang and Pei Chen. Delta-density based clustering with a divide-and-conquer strategy: 3dc clustering. *Pattern Recognition Letters*, 73:52–59, 2016.

[12] M. Lichman. UCI machine learning repository, 2013.

[13] Woongkee Loh and Hwanjo Yu. Fast density-based clustering through dataset partition using graphics processing units. *Information Sciences*, 308:94–112, 2015.

[14] James MacQueen et al. Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 281–297. Oakland, CA, USA., 1967.

[15] Satyasai Jagannath Nanda and Ganapati Panda. Design of computationally efficient density-based clustering algorithms. *Data and Knowledge Engineering*, 95:23–38, 2015.

[16] Mark EJ Newman. Modularity and community structure in networks. *Proceedings of the national academy of sciences*, 103(23):8577–8582, 2006.

[17] Truls A. Rklund, Michaela Tz, Johannes Gehrke, and Nils Grimsmo. Workload-aware indexing for keyword search in social networks. In *ACM CIKM*, pages 535–544, 2011.

[18] Alex Rodriguez and Alessandro Laio. Clustering by fast search and find of density peaks. *Science*, 344(6191):1492–1496, 2014.

[19] Alexander Strehl and Joydeep Ghosh. Cluster ensembles — a knowledge reuse framework for combining multiple partitions. *Journal of Machine Learning Research*, 3(3):583–617, 2002.

[20] Thanh N Tran, Ron Wehrens, and Lutgarde MC Buydens. Knn-kernel density-based clustering for high-dimensional multivariate data. *Computational Statistics & Data Analysis*, 51(2):513–525, 2006.

[21] Chang Dong Wang, Jian Huang Lai, and Jun Yong Zhu. Graph-based multiprototype competitive learning and its applications. *IEEE Transactions on Systems Man & Cybernetics Part C*, 42(6):934–946, 2012.

[22] Lijun Wang, Ming Dong, and Alexander Kotov. Multi-level approximate spectral clustering. In *IEEE ICDM*, pages 439–448, 2015.

[23] Jianhua Yin and Jianyong Wang. A dirichlet multinomial mixture model-based approach for short text clustering. In *ACM SIGKDD*, pages 233–242, 2014.

[24] Hongyan Zhang, Han Zhai, Liangpei Zhang, and Pingxiang Li. Spectral–spatial sparse subspace clustering for hyperspectral remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing*, 54(6):3672–3684, 2016.